

Intention Recognition, Commitment and the Evolution of Cooperation

The Anh Han*, Luís Moniz Pereira*, Francisco C. Santos*

* Centro de Inteligência Artificial (CENTRIA)

Departamento de Informática, Faculdade de Ciências e Tecnologia

Universidade Nova de Lisboa, 2829-516 Caparica, Portugal

Emails: h.anh@campus.fct.unl.pt, lmp@fct.unl.pt, fcsantos@fct.unl.pt

Abstract—Individuals make commitments towards others in order to influence others to behave in certain ways. Most commitments may depend on some incentive that is required to ensure that the action is in the agent's best interest and thus, should be carried out to avoid eventual penalties. Similarly, individuals may ground their decision on an accurate assessment of the intentions of others. Hence, both commitments and intention recognition go side by side in behavioral evolution. Here, we analyze the role played by the co-evolution of intention recognition plus the emergence of commitments, in the framework of the evolution of cooperative behavior. We resort to tools of evolutionary game theory in finite populations, showing how the combination of these two aspects of human behavior can enhance the emergent fraction of cooperative acts under a broad spectrum of configurations.

I. INTRODUCTION

Intention recognition (also called intention reading or understanding) is ubiquitous in many kinds of human interactions and communications, with much documented experimental evidence [1], [2], [3], [4]. Technically, intention recognition can be defined as a process of becoming aware of the intentions (or goals) of another agent or a group of other agents, inferring them through observed actions or effects on the environment [5], [6], [7], [8], [9], [10]. Intention recognition or intention reading is so critical for human social functioning and the development of key human abilities, such as language and culture, that it might have been shaped by natural selection [1], [4], [11], [12], [13].

But there are cases where it is difficult, if not impossible, to recognize the intentions of another agent. It might be your first interaction with someone in your life, and you have no information about him/her which can be used for intention recognition. You also might know someone well, but you still might have very little relevant information in a given situation to predict the intentions with high enough confidence. In such cases, the strategy of proposing a commitment, or manifesting an intention, can help to impose or clarify intentions of others¹.

Moreover, agents make commitments towards others when they give up options in order to influence others. Most com-

mitments depend on some incentive that is necessary to ensure that an action (or even an intention) is in the agent's interest and thus will be carried out in the future [17]. As previously, the capacity for using commitment strategies effectively is so important that natural selection may have shaped specialized capacities to make this possible [18], [19], [20], [21], [22], [23], [24], [25].

One of the commitments we all know is marriage. By giving up the option to leave someone else, spouses gain security and an opportunity for a much deeper relationship that would be impossible otherwise [23], [22], as it might be risky to assume a partner's intention of staying faithful without the commitment of marriage. A contract is another popular kind of commitment, e.g. for an apartment lease [22]. When it is risky to assume another agent's intention of being cooperative, arranging an appropriate contract provides incentives for cooperation. However, for example in accommodation rental, a contract is not necessary when the cooperative intention is of high certainty, e.g. when the business affair is between close friends or relatives. It said, arranging a commitment deal can be useful to encourage cooperation whenever intention recognition is difficult, or cannot be performed with sufficiently high certainty. On the other hand, arranging commitments is not free, and requires a specific capacity to set it up within a reasonable cost (for the agent to actually benefit from it) [23], [26] — therefore it should be avoided when opportune. In the case of marriage, partners sometimes choose to stay together without an official commitment when it might be too costly (e.g., it could be against parents' or families' wish, or it may need to be in secret because of their jobs) and/or they strongly trust each other's faithfulness (e.g., because of emotional attachment [27], [22]). In short, a combination of the two strategies, those of commitment and of intention recognition, seems unavoidable. Nevertheless, intention recognition without actual commitment can be enhanced by costly engagement gifts, in support of sexual selection and attachment [28], [29]. Furthermore, social emotions can act as ersatz commitment [27].

Here, we start from a simple model [30] of commitment formation, characterized by two key parameters: a punishment cost of failing commitment imposed on either side of a commitment deal, and the cost of managing it. It has been shown that, if a strong enough commitment deal can be

¹Intention is choice with commitment [14], [15], [16]. Once an agent intends to do something, it must settle on some state of affairs for which to aim, because of its resource limitation and in order to coordinate its future actions. Deciding what to do established a form of commitment [14], [16]. Proposing a commitment deal to another agent consists in asking it to express or clarify its intentions.

arranged, that is, with a small enough management cost and a large enough punishment cost, cooperation can emerge in a population of selfish agents.

On top of that model, using the tools of Evolutionary Game Theory (EGT) [31], [32], [33], [34], we show that combining intention recognition and commitment strategies in a reasonable way can lead to the emergence of improved cooperation, not able to be achieved solely by either strategy. Our study seeks what is a reasonable combination of commitment and intention recognition.

We shall do so in the setting of the Prisoner’s Dilemma (PD), a well-known game-theoretical framework to study the evolution of cooperation within populations of self-interested agents [35], [32], [36], [37]². In an interaction, each player has two options – to cooperate (C) or to defect (D) –, whereas defection is the dominant option, as it is always better to defect in one-shot interaction. Both players, if rational, will choose to defect, while they would be better off seeking mutual cooperation instead, thus leading to both the decrease of social welfare and individuals’ fitness.

It will be seen from our model that, in most of the cases, there is a wide range of combination of the intention recognition and commitment strategies, which leads to a strategy that performs better than either strategy solely – in the sense that the population spends more time in the homogeneous state of agents using that strategy [38], [39]. Our results suggest that, if one can recognize intentions of others with high enough confidence or certainty, one should rely more on it, especially when it is difficult to reach to a conceivably strong commitment deal. It helps to avoid the unnecessary cost of arranging and managing the deal. That is, in a transparent world where people have nothing to hide from each other, contracts are unnecessary. On the other hand, when intention recognition with high precision is difficult (due to, e.g. environment noise, agents have great incentives to hide intentions, or there is not enough observed actions), one should rely more on the commitment strategy, particularly if a reasonable deal can be envisaged.

The remainder of the paper is structured as follows. In Section II, we introduce the model of commitment, on top of which we integrate the co-evolution of commitment and intention recognition. In this section, we also describe the EGT methods to be used for analyzing our model. In Section III, we provide analytical and computer simulations obtained from our model. In Section IV, some discussions on the implication of the results are provided. The paper ends by proffering some concluding remarks.

II. MODELS AND METHODS

A. Models

We first summarize the commitment variant of the Prisoner’s Dilemma, then we describe our model for the combination of intention recognition and commitment.

²There are other social dilemmas such as the Stag Hunt and the Chicken Game [37], but the Prisoner’s Dilemma is known to represent one of the most difficult or fierce environments for cooperation to emerge.

1) *Commitment variant of the Prisoner’s Dilemma:* Let us consider a commitment variant of the Prisoner’s Dilemma game in which a new type of cooperator (denoted by COM_C) that, before each interaction, asks the co-player whether it commits to cooperate. If the co-player does not so commit, there is no interaction. Both players get 0. Otherwise, if the co-player commits, they then go on to play with each other in the present interaction. If the co-player keeps to its commitment, both players obtain the reward payoff, R . Note that here we do not yet take into account execution noise (see, e.g., [37]), i.e. the agents might mis-implement their intended choice, from cooperate to defect or vice versa. Thus, COM_C will never mis-implement the intended commitment choice, all the more so because commitment always entails an initial cost, thus being no point in proposing commitment when not intending to honor it. The payoffs of the commitment PD game, as we shall see, would make such bluffing players inevitably worse off. Otherwise (if the co-player fails its commitment), the proposing or focal player obtains the sucker payoff, S , and its co-player obtains the temptation payoff, T . However, the one that fails the commitment, whatever the player, will suffer a penalty cost for its non-defecting co-player to get a compensation. For simplicity, it is assumed that these two amounts are equal, denoted by δ . This cost can be a real monetary one, e.g., in the form of prior debit (e.g., in the case of accommodation rental) or of a punishment cost (e.g., commitment was performed in terms of a legal contract, and the one who fails commitment must pay a penalty cost), or an imaginary abstract value, e.g., public spread of good/bad reputation (bad reputation for the one that fails, and sympathy for the other), or even emotional suffering [23], [17], [40], [21]. How this cost is set up depends on the types of commitment at work, or the reason for which the commitment is believed to become fulfilled [41], [40], [30].

Two players that defect in an interaction obtain the punishment payoff, P . As usual, for the Prisoner’s Dilemma, the payoff entries satisfy the ordering, $T > R > P > S$, whereas the four possible outcomes can be written down as a payoff matrix

$$\begin{array}{cc} & C & D \\ \begin{array}{c} C \\ D \end{array} & \left(\begin{array}{cc} R, R & S, T \\ T, S & P, P \end{array} \right) \end{array}$$

For setting up a commitment, the proposer must pay a small cost, ϵ . The cost of proposing and setting up the commitment might be high, but it is reasonable to assume that this cost is very small compared to the mutual benefit of a cooperation strategy guaranteeing commitment, $\epsilon \ll R$.

We consider a finite population of constant size, consisting of four strategies: COM_C (as described above), C (always cooperates, without proposing a commitment deal), D (always defects, and does not commit when being asked to), and D_COM (always defects, though commits when being asked to). In [30], it is shown analytically and by computer simulations that COM_C dominates the population if the punishment cost δ is large enough compared to the management cost ϵ ,

thereby leading to the emergence of cooperation.

In each round, two random players are chosen from the population for an interaction. For the row player, the (average) payoff matrix reads

$$M_1 = \begin{matrix} & \begin{matrix} COM_C & C & D & D_COM \end{matrix} \\ \begin{matrix} COM_C \\ C \\ D \\ D_COM \end{matrix} & \begin{pmatrix} R - \epsilon/2 & R - \epsilon & -\epsilon & S + \delta - \epsilon \\ R & R & S & S \\ 0 & T & P & P \\ T - \delta & T & P & P \end{pmatrix} \end{matrix} \quad (1)$$

2) Combination of intention recognition and commitment:

We provide a new strategy, IRCOM, which combines the two strategies, those of intention recognition and commitment. In an interaction, IRCOM recognizes the intention (cooperates or defects) of its co-player [12]. A confidence level, cl , is assigned to the recognition result. It defines the degree of confidence (here in terms of probability) that IRCOM predicts the co-player's intention correctly³. In general, cl follows some probability distribution. As in a real intention recognition problem, the distribution should depend on the intention recognition method at work (how efficient it is), the environment IRCOM lives in (is it supportive for gathering relevant information for the recognition process, e.g. observability of co-players' direct and indirect interactions, perception noise, population structure), etc. For example, we can consider different distributions satisfying that the longer IRCOM survives, the more precisely or confidently it performs intention recognition; or, considering the repeated interaction setting in the framework of the iterated PD, the more IRCOM interacts with its co-player, the better it can recognize the co-player's intention (see intention recognition models for the iterated PD in [13], [12]).

We model cl by a continuous random variable X with probability density function $f(x, U)$, where U is a vector characterizing the factors that might influence cl , including the efficiency of the intention recognition model at work, the environmental factors (e.g., noise, population structure), and the interaction setting (repeated, one-shot, etc.).

If IRCOM is confident enough about the intention recognition process and result, that is cl is greater than a given, so-called, *confidence threshold* $\theta \in [0, 1]$, then in the current interaction IRCOM cooperates if the recognized intention of the co-player is to cooperate, and defects otherwise. The prediction is wrong with probability $(1-cl)$. For simplicity, we assume that the prediction is a (continuous) random variable, Y , uniformly distributed in $[0, 1]$. Hence, the probability that IRCOM utilizes intention recognition, but with an incorrect and correct prediction, respectively, can be written as joint

³In AI, the problem of intention recognition has been paid attention for several decades, and the main stream is that of probabilistic approaches [6], [7], [42], [43], [44]. They tackle the problem by assigning probabilities to conceivable intentions (conditional on the current observations), based on which the intentions are ranked. Similarly to [43], [44], [45], in our model, a degree of confidence in terms of a probability measure, is assigned to intentions.

probability distributions [46, Chapter 1] [47]

$$p_{ic} = P(X > \theta, Y < 1 - X) = \int_{\theta}^{+\infty} \int_{-\infty}^{1-x} f(x, U) dy dx, \quad (2)$$

$$p_c = P(X > \theta, Y > 1 - X) = \int_{\theta}^{+\infty} \int_{1-x}^{+\infty} f(x, U) dy dx. \quad (3)$$

If $cl \leq \theta$, i.e. IRCOM is not confident enough about its intention prediction, it behaves the same as COM_C (see above). The greater θ is, the more cautious IRCOM is about its intention recognition result. Obviously, if $\theta = 1$, IRCOM behaves identically to COM_C; and if $\theta = 0$, IRCOM behaves identically to a (pure) intention recognizer [12], [13] (see Figure 1).

We now replace COM_C with IRCOM, considering a population of four strategies, IRCOM, C, D, and D_COM. For the row player, the (average) payoff matrix reads

$$M = \theta M_1 + M_2 \quad (4)$$

where M_2 is the payoff matrix when IRCOM utilizes the intention recognition strategy, i.e. in the case $cl > \theta$. To derive M_2 , we consider the case that cl has a uniform distribution in the interval $[0, 1]$, i.e. $f(x, U) = 1$ for $x \in [0, 1]$ and 0 otherwise. Note that, on average, this can be considered as the distribution of a very inefficient intention recognition model because the confidence level or precision is a random number uniformly drawn from $[0, 1]$. As the prediction of the co-player's intention is only between two options, cooperate and defect, a random choice prediction already has a confidence level of 0.5.

Computing the integrals in Eqs. (2) and (3), we obtain: $p_{ic} = (1 - \theta) \frac{1-\theta}{2}$ and $p_c = (1 - \theta) \frac{1+\theta}{2}$. Hence,

$$M_2 = (1 - \theta) \begin{pmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & R & S & S \\ m_{31} & T & P & P \\ m_{41} & T & P & P \end{pmatrix} \quad (5)$$

where

$$\begin{aligned} m_{11} &= \frac{1}{4} [P(1 - \theta)^2 + (S + T)(1 - \theta)(1 + \theta) + R(1 + \theta)^2]; \\ m_{12} &= \frac{1}{2} [T(1 - \theta) + R(1 + \theta)]; \\ m_{13} &= \frac{1}{2} [S(1 - \theta) + P(1 + \theta)]; \\ m_{14} &= \frac{1}{2} [S(1 - \theta) + P(1 + \theta)]; \\ m_{21} &= \frac{1}{2} [S(1 - \theta) + R(1 + \theta)]; \\ m_{31} &= \frac{1}{2} [T(1 - \theta) + P(1 + \theta)]; \\ m_{41} &= \frac{1}{2} [T(1 - \theta) + P(1 + \theta)]. \end{aligned}$$

The main subject of the following analysis is to address, given the payoff entries of the PD, and the parameters of the commitment deal IRCOM can manage, how confident about the intention recognition result IRCOM should be in order to make a decision, without relying on the commitment proposing strategy. That is, if there is an optimal value of θ for an IRCOM to gain greatest net benefit.

B. Methods

Our analysis will be based on evolutionary game theory methods for finite populations [48], [38]. In the context of evolutionary game theory, the individuals' or agents' payoff represents their *fitness* or social *success*. The dynamics of

strategy change in a population is governed by social learning, that is, the most successful agents will tend to be imitated by the others. There are many ways to model social learning [32], [37], [49]. Adopting one of the most frequently used ones, we shall consider the so-called pairwise comparison rule [50], which assumes that an agent A with fitness f_A adopts the strategy of another agent B with fitness f_B with probability given by

$$\frac{1}{1 + e^{-\beta(f_B - f_A)}},$$

where β controls the ‘imitation strength’, i.e., how strongly the agents are basing the decision to imitate on fitness comparisons. For $\beta = 0$, we obtain the limit of neutral drift – the imitation decision is random. For large β , imitation becomes increasingly deterministic.

In the absence of mutations, the end states of evolution are inevitably monomorphic: once such a state is reached, imitation cannot produce any change. We thus further assume that, with a certain mutation probability $\mu > 0$ (also dubbed the exploration rate [51]), an agent switches randomly to a different strategy without imitating any other agent. The resulting Markov Chain has a stationary distribution, which characterizes the average time the population spends in each of these monomorphic end states. Yet, for arbitrary exploration rates and number of strategies, stationary distributions are often cumbersome to compute [39], [34], [52].

Fortunately, in the case of small exploration or mutation rates, analytical computation of this stationary distribution can conveniently be computed [53], [38], [39], [20]. The small exploration rates guarantee that any newly occurred mutant in a homogeneous population will fixate or become extinct long before the occurrence of another mutation. Hence, the population will always consist of at most two strategies in co-presence. This allows one to describe the evolutionary dynamics of our population in terms of a reduced Markov Chain, whose size is equal to the number of strategies being considered (which is 4 in our case), and each state represents a possible monomorphic end state of the population associated with one of the strategies. The transitions between states are defined by the fixation probabilities of a single mutant of one strategy in a homogeneous population of individuals adopting another strategy.

More precisely, let N be the size of the population. Suppose there are at most two strategies in the population, say, k agents using strategy A ($0 \leq k \leq N$) and $(N - k)$ agents using strategy B. Thus, the (average) payoff of the agent that uses A or uses B can be written as follows, respectively,

$$\begin{aligned} \Pi_A(k) &= \frac{(k-1)\pi_{A,A} + (N-k)\pi_{A,B}}{N-1} \\ \Pi_B(k) &= \frac{k\pi_{B,A} + (N-k-1)\pi_{B,B}}{N-1} \end{aligned} \quad (6)$$

where $\pi_{X,Y}$ stands for the payoff an agent using strategy X obtained in an interaction with another agent using strategy Y , given by the payoff matrix (4).

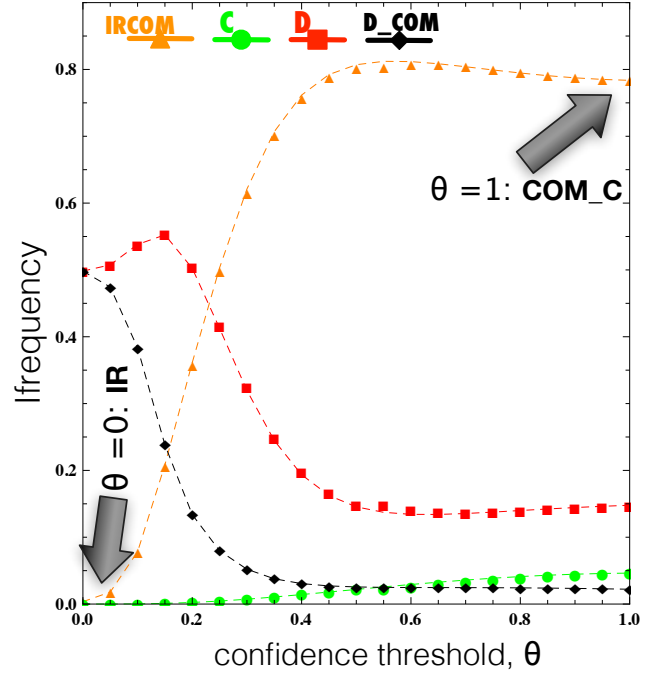


Figure 1: Frequency of each strategy as a function of confidence threshold, θ . Symbols indicate results from computer simulations (averaged over 10^7 interactions for each pair of strategies), and dashed curves show the exact numerical results (see Methods). In a population of IRCOM, COM_D, C, and D individuals, for large enough θ , the population spends most of the time in the homogeneous state of IRCOM. The performance of IRCOM decreases when θ is too high. It implies that IRCOM should not be too cautious about its intention recognition capacity, i.e. not be too careful to always propose commitment instead of believing in its intention recognition result. Parameters: $\delta = 4$; $\epsilon = 0.05$; payoff entries, $T = 2$, $R = 1$, $P = 0$, $S = -1$; population size, $N = 100$; imitation strength, $\beta = 0.1$.

Now, the probability to change, by ± 1 , the number k of agents using strategy A at each time step can be written as

$$T^\pm(k) = \frac{N-k}{N} \frac{k}{N} \frac{1}{1 + e^{\mp\beta[\Pi_A(k) - \Pi_B(k)]}}. \quad (7)$$

The fixation probability of a single mutant with a strategy A in a population of $(N - 1)$ agents using B is given by [50], [54], [53], [38], [39]

$$\rho_{B,A} = \frac{1}{1 + \sum_{i=1}^{N-1} \prod_{j=1}^i \frac{T^-(j)}{T^+(j)}}. \quad (8)$$

In the limit of neutral selection ($\beta = 0$), $T^-(j) = T^+(j) \forall j$. Thus, $\rho_{B,A} = 1/N$. Considering a set $\{1, \dots, q\}$ of different strategies, these fixation probabilities determine a transition matrix $M = \{T_{ij}\}_{i,j=1}^q$, with $T_{ij,j \neq i} = \rho_{ji}/(q-1)$ and $T_{ii} = 1 - \sum_{j=1, j \neq i}^q T_{ij}$, of a Markov Chain. The normalized eigenvector associated with the eigenvalue 1 of the transposed of M provides the stationary distribution described above [54], [53], [38], [39], describing the relative time the population spends adopting each of the strategies.




Figure 2: Frequency of IRCOM as a function of confidence threshold, θ . (a) We plot for different values of management cost, ϵ . In a population of IRCOM, COM_D, C, and D individuals, for large enough θ and small enough ϵ , the population spends most of the time in the homogeneous state of IRCOM. The smaller ϵ , the better the performance of IRCOM. The performance of IRCOM decreases when θ is too high, and the greater ϵ , the more significant the decrease is. It implies that the more costly the management of the commitment deal, the more beneficial it is to rely on intention recognition. (b) We plot for different values of punishment cost, δ . In a population of IRCOM, COM_D, C, and D individuals, for large enough θ and large enough δ , the population spends most of the time in the homogeneous state of IRCOM. The greater δ , the better the performance of IRCOM. The performance of IRCOM decreases when θ is too high, and the decrease is more significant when δ is smaller. It implies that the weaker the commitment deal can be arranged, the more beneficial it is to rely on intention recognition alone more often. Parameters: $\delta = 4$ in panel (a) and $\epsilon = 0.05$ in panel (b); payoff entries, $T = 2$, $R = 1$, $P = 0$, $S = -1$; population size, $N = 100$; imitation strength, $\beta = 0.1$. The results are computed numerically (see Methods).

III. RESULTS

To start with, we compute the stationary distributions analytically and resorting to agent-based simulations (see Methods and Figure 1). The results show that, for a large range of θ , IRCOM performs better than COM_C (i.e. IRCOM with $\theta = 1$), whereas the population spends most of the time in the homogenous state of IRCOM. However, when the confidence threshold is low, defection becomes dominant. This said, when the intention recognition is not of high enough certainty—that is, IRCOM is not confident enough about whether its co-player intends to cooperate or to defect in the current interaction—it would be better off counting on the commitment strategy (this also can be observed in several different configurations in Figure 2).

In Figure 2, we analyze the influence of the strength of the commitment deal which can be arranged, on how the intention recognition and commitment strategies can be combined appropriately. Note that the greater the punishment cost, δ , and the smaller the management cost, ϵ , the stronger the commitment deal. Namely, in Figure 2a, fixing δ , we plot for different values of management cost, ϵ . The performance of IRCOM decreases when θ is too high, and the decrease is more dramatic when ϵ is greater. It means that the costlier the

management of the commitment deal, the more beneficial it is to rely on intention recognition. Next, in Figure 2b, fixing ϵ , we plot for different values of punishment cost, δ . The performance of IRCOM decreases when θ is too high, and the decrease is more dramatic when δ is smaller. In short, these results imply that the weaker the commitment deal can be arranged, the more beneficial it is to rely on intention recognition.

So far the model has been studied with respect to a very inefficient intention recognition model, where cl is uniformly distributed in $[0, 1]$. It is not surprising that the performance of the intention recognition strategy solely—which corresponds to IRCOM with $\theta = 0$ (see Figure 1)—is very poor (Figures 1 and 2). In the sequel, let us study the model using more efficient intention recognition models.

We consider different probability distributions of the confidence level cl , reflecting different levels of efficiency or precision of the intention recognition model at work, given the relevant factors (noise, environment factors, interaction settings, etc.) (Figure 3). Namely, in an increasing order of efficiency, cl is uniformly drawn from intervals $[0, 1]$, $[0.5, 1]$,

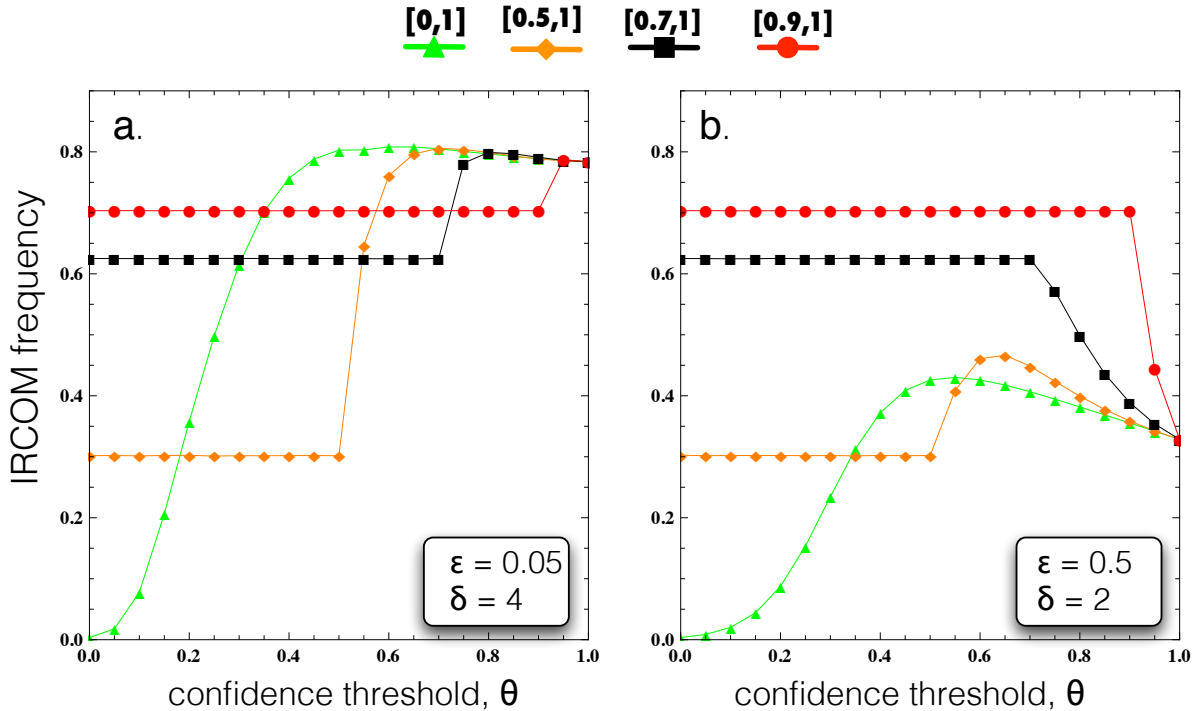


Figure 3: Frequency of IRCOM as a function of confidence threshold, θ , in a population of IRCOM, COM_D, C, and D individuals. We consider different probability distributions of the confidence level cl , reflecting the efficiency or precision of the intention recognition model at work, given the relevant factors (noise, environment, etc.). Namely, cl is uniformly drawn from $[0, 1]$, $[0.5, 1]$, $[0.7, 1]$, and $[0.9, 1]$. The results show that if a strong commitment deal can be arranged (panel a), it is better to rely on the commitment strategy using a high enough confidence threshold—even when the intention recognition model is very efficient, while it is more beneficial to rely, even exclusively, on the intention recognition strategy if it is efficient enough, in the case that only weak commitment deals can be arranged (panel b). Parameters: $\epsilon = 0.05$, $\delta = 4$ (panel a) and $\epsilon = 0.5$, $\delta = 2$ (panel b); payoff entries, $T = 2$, $R = 1$, $P = 0$, $S = -1$; population size, $N = 100$; imitation strength, $\beta = 0.1$; The payoff matrices in all cases are derived by averaging 10^7 interactions of each pair of strategies. The results are computed numerically (see Methods).

$[0.7, 1]$, and $[0.9, 1]$ ⁴. The results show that, whenever the intention recognition model is efficient enough, the intention recognition strategy solely (i.e. IRCOM with $\theta = 0$) performs quite well, complying with the results obtained in [12], [13], where concrete intention recognition models are deployed.

However, when a quite strong commitment deal can be envisaged (Figure 3a), arranging it can still glean some evolutionary advantage. But in case that only weak commitment deals can be arranged (Figure 3b), it is then more beneficial to rely, even exclusively, on the intention recognition strategy should it be efficient enough.

IV. DISCUSSION

A general implication of our analysis is that an appropriate combination of the two strategies of commitment and intention

⁴In the context of iterated interactions (e.g. in the framework of the iterated Prisoner's Dilemma), these levels of efficiency can be achieved (on average) by considering large enough numbers of interactions between two players (or high enough probabilities of a next interaction or 'the shadow of future' [36], [37]), given that the noise is small enough. Normally, the more an intention recognizer interacts with a fixed co-player, the better it predicts its co-player's intention. For example, this holds for the two intention recognition models described in [12], [13]. Furthermore, in [55], the authors present experimental evidence showing that, in a one-shot PD, subjects of only brief acquaintance were able to recognize players with an intention to defect with more than twice chance accuracy.

recognition often leads to a strategy that performs better than either one solely. It is advantageous to rely on the intention recognition strategy (when reaching sufficiently high confidence about its result) because it helps to avoid the cost of arranging and managing commitment deals, especially when no strong deals can be arranged or envisaged. This result has a similar implication to that obtained in [56], where the authors show that overconfidence might give evolutionary advantage to its holders. In our model, an IRCOM can gain extra net benefit if it is a little overconfident (that is, when using sufficiently small θ), taking risk to rely on intention recognition result instead of arranging some commitment deal. Differently, because in our model IRCOM is further guaranteed by an efficient strategy of commitment, being over-overconfident (that is, using too small θ) and relying exclusively on intention recognition might prevent it from opportunely gaining benefit from the commitment strategy—especially in case the intention recognition model at work is not efficient. It said, the performance of overconfident individuals [56] can be enhanced by relying on the commitment strategy when they need to muster overly high courage (say, in order to decide to claim some resource).

In the framework where intention recognition is difficult and of high risk, for example, climate change negotiation [57],

[58], [52], military setting (comprising a lot of bluffing) [41], [59], and international relationships [60], our model suggests arranging a strong commitment deal.

V. CONCLUSION

Assume simply that we are given an intention recognition method, that affords us a degree of confidence distribution cl about its predictions, with regard to the intentions of others, and hence their future actions, typically on the basis of their seen actions and surrounding historical and present circumstances. Assume too some commitment model is given us about providing mutual assurances, and involving an initial cost and a penalty for defaulting.

We have shown how to combine together one such general intention recognition method, with a specific commitment model defined for playing the Prisoner's Dilemma (PD), in the setting of Evolutionary Game Theory (EGT), by means of a single payoff matrix extended with a new kind of player, IRCOM, which chooses whether to go by the result of its intention recognition method about a co-player's next move, or to play by the commitment strategy, depending on whether its level of confidence on the intention prediction cl exceeds or not some a given confidence threshold θ . Our results indicate that IRCOM is selected by evolution for a broad range of parameters and confidence thresholds.

Then we have studied, for a variety of cl and θ , in the context of PD in EGT, how IRCOM performs in the presence of other well-known non-committing strategies (always cooperate, C, and always defect, D) – plus the strategy that commits when being asked to, but always defects, D_COM. Analytical and simulation results show under which circumstances, for different cl and θ , and distinct management and punishment costs, ϵ and δ , that the new combined strategy IRCOM proves advantageous and to what degree. And indeed IRCOM proves to be adaptably advantageous over those other strategies, in all circumstances, for a quite small confidence level onwards.

Much remains to be done with respect to further consideration of combining the two strategies of intention recognition and commitment. The two go often together, and not just in the basic way we have examined. Indeed they are the two sides of the same coin, one side being an attempt to identify an intention, the other side being a manifestation of an intention. For one, we only considered the case where intention recognition comes first in order to decide on a commitment proposal. But in general, once a commitment is made, intention recognition is a paramount method to follow up on whether the commitment will be honored, on the basis of detecting or not the intermediate actions leading up to commitment fulfillment. Social organizations rely on these mechanisms to structure themselves. Furthermore, the information about commitments can be used to enhance intention recognition.

It seems to us that intention recognition, and its use in the scope of commitment, is a foundational cornerstone where we should begin at, naturally followed by the capacity to establish

and honor commitments, as a tool towards the successive construction of collective intentions and social organization [61], [62]. Finally, one hopes that understanding these capabilities can be useful in the design of efficient self-organized and distributed engineering applications [63], from bio and socio-inspired computational algorithms to swarms of autonomous robotic agents.

ACKNOWLEDGMENT

HTA and FCS acknowledge the support from FCT-Portugal (grant SFRH/BD/62373/2009 and R&D project PTDC/FIS/101248/2008, respectively).

REFERENCES

- [1] A. L. Woodward, J. A. Sommerville, S. Gerson, A. M. Henderson, and J. Buresh, "The emergence of intention attribution in infancy," in *The Psychology of Learning and Motivation*, ser. Psychology of Learning and Motivation, B. H. Ross, Ed. Academic Press, 2009, vol. 51, pp. 187 – 222.
- [2] A. N. Meltzoff, "The framework for recognizing and becoming an intentional agent," *Acta Psychol Amst*, vol. 124, no. 1, pp. 26–43, 2007.
- [3] M. Tomasello, *The Cultural Origins of Human Cognition*. Harvard University Press, 1999.
- [4] —, *Origins of Human Communication*. MIT Press, 2008.
- [5] H. A. Kautz and J. F. Allen, "Generalized plan recognition," in *Proceedings of 1986 Conference of the American Association for Artificial Intelligence (AAAI'1986)*. AAAI, 1986, pp. 32–37.
- [6] E. Charniak and R. P. Goldman, "A Bayesian model of plan recognition," *Artificial Intelligence*, vol. 64, no. 1, pp. 53–79, 1993.
- [7] C. Heinze, "Modeling intention recognition for intelligent agent systems," Ph.D. dissertation, The University of Melbourne, Australia, 2003.
- [8] L. M. Pereira and T. A. Han, "Intention recognition via causal Bayes networks plus plan generation," in *Progress in Artificial Intelligence, Proceedings of 14th Portuguese International Conference on Artificial Intelligence (EPIA'09)*. Springer LNAI 5816, October 2009, pp. 138–149.
- [9] —, "Intention recognition with evolution prospection and causal Bayesian networks," in *Computational Intelligence for Engineering Systems 3: Emergent Applications*. Springer, 2011, pp. 1–33.
- [10] T. A. Han and L. M. Pereira, "Intention-based decision making with evolution prospection," in *Proceedings of 15th Portuguese Conference on Artificial Intelligence (EPIA'2011)*, L. Antunes and H. Pinto, Eds. LNAI 7026, Springer LNAI, 2011, pp. 254–267.
- [11] B. Skyrms, *Evolution of the Social Contract*. Cambridge University Press, 1996.
- [12] T. A. Han, L. M. Pereira, and F. C. Santos, "Intention recognition promotes the emergence of cooperation," *Adaptive Behavior*, vol. 19, no. 3, pp. 264–279, 2011.
- [13] —, "The role of intention recognition in the evolution of cooperative behavior," in *Proceedings of the 22nd international joint conference on Artificial intelligence (IJCAI'2011)*, T. Walsh, Ed. AAAI, 2011, pp. 1684–1689.
- [14] P. R. Cohen and H. J. Levesque, "Intention is Choice with Commitment," *Artificial Intelligence*, vol. 42, no. 2-3, pp. 213–261, 1990.
- [15] M. E. Bratman, *Intention, Plans, and Practical Reason*. The David Hume Series, CSLI, 1987.
- [16] O. Roy, "Thinking before acting: Intentions, logic, rational choice," Ph.D. dissertation, ILLC Dissertation Series DS-2008-03, Amsterdam., 2009.
- [17] H. Gintis, "Beyond selfishness in modeling human behavior," in *Evolution and the capacity for commitment*, R. M. Nesse, Ed. New York: Russell Sage, 2001.
- [18] B. Skyrms, *Signals: Evolution, Learning, and Information*. Oxford University Press, 2010.
- [19] A. Robson, "Efficiency in evolutionary games: Darwin, Nash, and the secret handshake," *Journal of Theoretical Biology*, vol. 144, no. 3, pp. 379–396, 1990.
- [20] F. C. Santos, J. M. Pacheco, and B. Skyrms, "Co-evolution of pre-play signaling and cooperation," *Journal of Theoretical Biology*, vol. 274, no. 1, pp. 30–35, 2011.

- [21] M. Ruse, "Morality and commitment," in *Evolution and the capacity for commitment*, R. M. Nesse, Ed. New York: Russell Sage, 2001, pp. 221–236.
- [22] R. H. Frank, "Cooperation through Emotional Commitment," in *Evolution and the capacity for commitment*, R. M. Nesse, Ed. New York: Russell Sage, 2001, pp. 55–76.
- [23] R. M. Nesse, "Natural selection and the capacity for subjective commitment," in *Evolution and the capacity for commitment*, R. M. Nesse, Ed. New York: Russell Sage, 2001, pp. 1–44.
- [24] de Vos, R. Smaniotto, and D. Elsas, "Reciprocal altruism under conditions of partner selection," *Rationality and Society*, vol. 13, no. 2, pp. 139–183, 2001.
- [25] I. Back and A. Flache, "The Adaptive Rationality of Interpersonal Commitment," *Rationality and Society*, vol. 20, no. 1, pp. 65–83, 2008.
- [26] R. M. Nesse, *Evolution and the capacity for commitment*, ser. Russell Sage Foundation series on trust. Russell Sage, 2001.
- [27] R. H. Frank, *Passions Within Reason: The Strategic Role of the Emotions*. W. W. Norton and Company, 1988.
- [28] G. F. Miller and P. M. Todd, "Mate choice turns cognitive," *Trends in Cognitive Sciences*, vol. 2, no. 5, pp. 190–198, 1998.
- [29] M. G. Haselton and D. M. Buss, "Error management theory: A new perspective on biases in cross-sex mind reading," *Journal of Personality and Social Psychology*, vol. 78, no. 1, pp. 81–91, 2001.
- [30] T. A. Han, L. M. Pereira, and F. C. Santos, "The emergence of commitments and cooperation," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'2012)*, 2012, (Forthcoming). [Online]. Available: <http://centria.fct.unl.pt/Imp/publications/online-papers/aamas2012.pdf>
- [31] J. Maynard-Smith, *Evolution and the Theory of Games*. Cambridge: Cambridge University Press, 1982.
- [32] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.
- [33] M. A. Nowak, *Evolutionary Dynamics*. Harvard University Press, Cambridge, MA, 2006.
- [34] K. Sigmund, H. D. Silva, A. Traulsen, and C. Hauert, "Social learning promotes institutions for governing the commons," *Nature*, vol. 466, p. 7308, 2010.
- [35] R. Axelrod, *The Evolution of Cooperation*. Basic Books, ISBN 0-465-02122-2, 1984.
- [36] M. A. Nowak, "Five rules for the evolution of cooperation," *Science*, vol. 314, no. 5805, p. 1560, 2006, doi: 10.1126/science.1133755.
- [37] K. Sigmund, *The Calculus of Selfishness*. Princeton University Press, 2010.
- [38] L. A. Imhof, D. Fudenberg, and M. A. Nowak, "Evolutionary cycles of cooperation and defection," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, pp. 10797–10800, 2005.
- [39] C. Hauert, A. Traulsen, H. Brandt, M. A. Nowak, and K. Sigmund, "Via freedom to coercion: The emergence of costly punishment," *Science*, vol. 316, pp. 1905–1907, 2007.
- [40] J. Hirshleifer, "Game-theoretic interpretations of commitment," in *Evolution and the capacity for commitment*, R. M. Nesse, Ed. New York: Russell Sage, 2001, pp. 77–93.
- [41] T. C. Schelling, *The strategy of conflict*. London: Oxford University Press, 1990.
- [42] H. Bui, S. Venkatesh, and G. West, "Policy recognition in the abstract hidden markov model," *Journal of Artificial Intelligence Research*, vol. 17, pp. 451–499, 2002.
- [43] N. Blaylock and J. Allen, "Statistical goal parameter recognition," in *Proceedings of the 14th International Conference on Automated Planning and Scheduling (ICAPS'04)*, S. Zilberstein, J. Koehler, and S. Koenig, Eds. AAAI, 2004, pp. 297–304.
- [44] M. G. Armentano and A. Amandi, "Goal recognition with variable-order markov models," in *Proceedings of the 21st international joint conference on Artificial intelligence*, 2009, pp. 1635–1640.
- [45] T. A. Han and L. M. Pereira, "Context-dependent incremental intention recognition through Bayesian network model construction," in *Proceedings of the Eighth UAI Bayesian Modeling Applications Workshop (UAI-AW 2011)*, A. Nicholson, Ed., vol. 818. CEUR Workshop Proceedings, 2011, pp. 50–58. [Online]. Available: <http://ceur-ws.org/Vol-818/paper7.pdf>
- [46] A. Gut, *An Intermediate Course in Probability*, 2nd ed. Springer Publishing Company, Incorporated, 2009.
- [47] T. A. Han, A. Traulsen, and C. S. Gokhale, "On equilibrium properties of evolutionary multiplayer games with random payoff matrices," *Theoretical Population Biology*, 2012.
- [48] M. A. Nowak, A. Sasaki, C. Taylor, and D. Fudenberg, "Emergence of cooperation and evolutionary stability in finite populations," *Nature*, vol. 428, pp. 646–650, 2004.
- [49] L. Rendell, R. Boyd, D. Cownden, M. Enquist, K. Eriksson, M. W. Feldman, L. Fogarty, S. Ghirlanda, T. Lillicrap, and K. N. Laland, "Why copy others? insights from the social learning strategies tournament," *Science*, vol. 328, no. 5975, pp. 208–213, 2010. [Online]. Available: <http://www.sciencemag.org/content/328/5975/208.abstract>
- [50] A. Traulsen, M. A. Nowak, and J. M. Pacheco, "Stochastic dynamics of invasion and fixation," *Phys. Rev. E*, vol. 74, p. 11909, 2006.
- [51] A. Traulsen, C. Hauert, H. De Silva, M. A. Nowak, and K. Sigmund, "Exploration dynamics in evolutionary games," *Proc. Natl. Acad. Sci. USA*, vol. 106, no. 3, pp. 709–712, 2009.
- [52] F. C. Santos and J. M. Pacheco, "Risk of collective failure provides an escape from the tragedy of the commons," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, no. 26, pp. 10421–10425, 2011.
- [53] D. Fudenberg and L. A. Imhof, "Imitation processes with small mutations," *Journal of Economic Theory*, vol. 131, pp. 251–262, 2005.
- [54] S. Karlin and H. E. Taylor, *A First Course in Stochastic Processes*. Academic Press, New York, 1975.
- [55] R. H. Frank, T. Gilovich, and D. T. Regan, "The evolution of one-shot cooperation: An experiment," *Ethology and Sociobiology*, vol. 14, no. 4, pp. 247–256, 1993.
- [56] D. D. P. Johnson and J. H. Fowler, "The evolution of overconfidence," *Nature*, vol. 477, no. 7364, pp. 317–320, 2011.
- [57] N. Raihani and D. Aitken, "Uncertainty, rationality and cooperation in the context of climate change," *Climatic Change*, vol. 108, no. 1, pp. 47–55, 2011.
- [58] M. Milinski, D. Semmann, H. J. Krambeck, and J. Marotzke, "Stabilizing the Earth's climate is not a losing game: Supporting evidence from public goods experiments," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 103, pp. 3994–3998, 2006.
- [59] B. A. Leeds, "Alliance Reliability in Times of War: Explaining State Decisions to Violate Treaties," *International Organization*, vol. 57, no. 04, pp. 801–827, 2003. [Online]. Available: <http://journals.cambridge.org/action/displayAbstract?fromPage=online&aid=183071>
- [60] C. Lockhart, "Flexibility and commitment in international conflicts," *International Studies Quarterly*, vol. 22, no. 4, pp. 545–568, 1978.
- [61] J. R. Searle, *The Construction of Social Reality*. New York: The Free Press, 1995.
- [62] —, *Making the Social World: The Structure of Human Civilization*. Oxford University Press, 2010.
- [63] E. Bonabeau, M. Dorigo, and G. Theraulaz, *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, USA, 1999.